

---

# The RIACS Mail System

*Matt Bishop*

February 18, 1987

Research Institute for Advanced Computer Science  
NASA Ames Research Center

RIACS TR 87.6

(NASA-CR-187286) THE RIACS MAIL SYSTEM  
(Research Inst. for Advanced Computer  
Science) 26 p

N90-71405

00/81      Unclas  
            0295380

# RIACS

Research Institute for Advanced Computer Science

---

# **The RIACS Mail System**

*Matt Bishop*

Research Institute for Advanced Computer Science  
NASA Ames Research Center  
Moffett Field, CA 94035

## **ABSTRACT**

This document describes the configuration of the RIACS mail system, and how the mail configuration files implement this design. The rationale behind the design, as well as the configuration files, is discussed. This paper concludes with a description of features that may be used to debug the configuration files, as well as some basic guidelines on how to trace problems with configuration files that have been installed.

February 18, 1987

# The RIACS Mail System

Matt Bishop

Research Institute for Advanced Computer Science  
NASA Ames Research Center  
Moffett Field, CA 94035

## 1. Introduction

Users who send or receive letters over an electronic network deal with only a small part of the mail system. Typically, sending mail involves several levels of services and protocols, summarized in Figure 1. (This model is called *the MHS model* and is from ISO X400.)




Figure 1. The MHS model.

In this figure, the person who wants to send the letter is called the *originator*. He uses a *user agent*, which is a program that enables him to type and format the letter. The user agent then hands the letter to a *message transport system*. This system is composed of one or more *message transport agents*; each agent accepts the message, determines if the recipient is on its machine, and if so passes the message to the appropriate user agent. If not, it determines which machine it should deliver the message to, and passes the letter on to the message transport agent on that machine. Finally, the letter reaches a user agent on the recipient's machine, and the user agent enables the recipient to read and possibly reply to the letter.

RIACS has many user agents available, ranging from the simple *binmail*(1) [UPM84] to the highly sophisticated *MH Mail Handling System* [ROSE86]. This manual leaves the description of the user agents to other documents; it instead concentrates on the message transport agents.

The next section describes the configuration of the machines at RIACS; the sections after it go line-by-line through the gateway and non-gateway *sendmail* configuration files. The final sections discuss how to use *sendmail* [ALLM84a, ALLM84b] to debug a configuration file, and how to install new configuration files. These last sections assume the reader is familiar enough with configuration files that the notation and general layout need not be reviewed; if not familiar with these, the reader should have a copy of [ALLM84b] readily available when reading these sections.

## 2. RIACS Mail Configuration

The RIACS environment is quite heterogeneous; it consists of a VAX-11/730† which serves as a gateway, a Sequent Balance 21000, an Intel Hypercube, a Ridge 32/V, a Silicon Graphics IRIS 3500, and several SUN 3's, all running some version of UNIX,‡ some sharing file systems

†VAX is a Trademark of Digital Equipment Corporation.

‡UNIX is a Trademark of Bell Laboratories.

using NFS and others not. The following table shows the names of these machines:

Table 1. RIACS Computers and their Names	
<i>host name</i>	<i>computer</i>
icarus	VAX-11/730
hydra	Sequent Balance 21000
cube	Intel Hypercube
daedalus	Ridge 32/V
pegasus	Silicon Graphics IRIS 3500
lavalite	Sun 3 (has file system)
miranda	Sun 3 (has file system; server)
phun	Sun 3 (client sharing <i>miranda's</i> file system)
zeus	Sun 3 (client sharing <i>miranda's</i> file system)
clavier	Sun 3 (client sharing <i>miranda's</i> file system)
dora	Sun 3 (client sharing <i>miranda's</i> file system)

There is also an Evans and Sutherlin P300 graphics system, two Apple Macintoshes, and an IBM PC used for financial matters.

Of these machines, *icarus*, *hydra*, *daedalus*, and the Suns are the only ones which allow the full range of mailing facilities. *Cube* allows local mail to be sent, but does not have any facility for sending mail to another machine except via *uucp*. *Pegasus* does not allow any nonlocal mail to be sent; it, the Evans and Sutherlin, the Macintoshes, and the IBM PC do not interact with the mail system.

There are a number of configurations that allow mail to be sent to nonlocal or remote machines. One is to require each host to have all the information needed, so when a letter is sent to *csl.sri.com* from *hydra*, *hydra's* mail router will send the letter directly to *csl.sri.com*. On first blush, this seems to be the best option; however, it has drawbacks. Not every host can transmit all types of mail; for example, *icarus* is the only host that can be used to send *uucp* mail, because it is the only host with telephone lines connected. If some distant site changes its internet address, all machines must have their host tables updated (since some machines, such as *hydra*, do not use domain name resolving); worse, when a configuration file is changed to reflect a change in the way mail is handled, all configuration files must be changed; this leads to problems of mail being caught in transit when the change occurs.

Since all nonlocal mail must be sent through gateway machine, why not send all nonlocal mail there for forwarding? In this configuration, whenever any letter being mailed has an address on some other machine, it is sent directly to *icarus*. *Icarus* then decides how to send it to its destination. For example, if a letter is sent to *csl.sri.com* from *hydra*, *hydra's* mail router will send the letter to *icarus*; then *icarus's* mail router will send the letter on to *csl.sri.com*. This scheme has advantages over the previous one, in that all hosts except the gateway need only know which host is the gateway; if there is any mail on the host for any other machine, it is sent to the gateway, which then decides how best to deliver it. If some distant site changes its internet address, only the host table on the gateway need be changed; the tables on the other hosts may be updated as convenient, but doing so (or failing to do so) will not affect the transfer of mail. One would suspect that this configuration is worse than the previous one, because if *icarus* goes down, no mail can be sent off site. However, since all our network traffic goes through *icarus*, if *icarus* is down, no off site mail could be sent anyway. The only problem is that inter-site mail will not go from one machine to another when *icarus* is down.

There is one other advantage to the latter configuration. RIACS uses the domain naming system of the ARPANET. As a result, having a central machine provides a very easy way to decide which machine is to be listed as "riacs.edu". All other machines are invisible to off site systems; any letter going out, regardless of the machine from which it originated, has the return address changed to read "*user@riacs.edu*". In essence, we can reconfigure our machines, change their internet addresses, and so forth without having to have people throughout the country update their host tables. Only the address of the machine answering to "riacs.edu" need be kept

the same.

For these reasons, RIACS uses the latter configuration. Pictorially, this system is arranged as shown in Figure 2.




Figure 2. The RIACS Mail Configuration

There are two different configuration files; the first, the *gateway* version, resides on the RIACS gateway and handles routing over a variety of communications media; the second, the *non-gateway* version, resides on all other hosts. This next section describes the *gateway* configuration file.

### 3. *Sendmail* Configuration File Preliminaries

This section describes options, macros, classes, mailers, and other characteristics of the configuration files. Rulesets are described in the next section.

#### 3.1. Macros

Certain macros and options are set in both the gateway and non-gateway configuration files. Some are necessary for *sendmail* to work; others are peculiar to RIACS. In this section, those common settings are explained.

The first three macros in the gateway's file set the domain name **N**, the uucp host name **U**, and the version **Z** of this *sendmail* configuration file; in the non-gateway file, the uucp host name is not given but the gateway host **G** is. *Sendmail* also requires six macros to be defined; otherwise, it will not function properly. Table 2 summarizes the definitions of these macros.

Table 2. RIACS Settings for Macros		
macro	gateway setting	non-gateway setting
<b>G</b>	<i>not set</i>	icarus.riacs.edu
<b>N</b>	riacs.edu	
<b>U</b>	riacs.uucp	<i>not set</i>
<b>Z</b>	2.0G	2.0N
<b>j</b>	\$w	
<b>e</b>	\$j Sendmail \$v/\$Z ready at \$b	
<b>l</b>	From \$g \$d	
<b>n</b>	MAILER-DAEMON	
<b>o</b>	.:%@!^=//	
<b>q</b>	\$?x\$x \$.<\$g>	

### 3.2. Options

The configuration files also set many options for *sendmail*; both the gateway and non-gateway configuration files set the same ones. Table 3 summarizes their settings.

Table 3. RIACS Configuration File Options			
option	RIACS setting	option	RIACS setting
<b>A</b>	/usr/lib/aliases	<b>a</b>	10
<b>B</b>	_	<b>d</b>	b
<b>F</b>	0600	<b>g</b>	1
<b>H</b>	/usr/lib/sendmail.hf	<b>o</b>	set
<b>L</b>	9	<b>r</b>	1h
<b>Q</b>	/usr/spool/mqueue	<b>s</b>	set
<b>S</b>	/usr/lib/sendmail.st	<b>t</b>	PST,PDT
<b>T</b>	3d	<b>u</b>	1
<b>W</b>	*	<b>x</b>	8
<b>X</b>	12		

Some of these options require a bit of explanation.

- A,a** These refer to the alias file feature of *sendmail*. The alias file is located in "/usr/lib/aliases", and if *sendmail* detects the alias file is being rebuilt it will wait up to ten minutes before deciding the rebuild has failed and initiate one of its own.
- B** Older mailers allowed the use of blank characters in addresses; this causes all sorts of problems on many systems. This option directs *sendmail* to replace all blanks with underscores "\_". RIACS used to use periods ".", but this poses problems in a domain system; an underscore does not.
- d** This option instructs the daemon to run in background mode; the daemon will disconnect from the terminal.
- F,Q** The value of these options is the mode and location of the temporary and queue files used by *sendmail*. It is a RIACS policy that the mode shall be at least mode 0660, and preferably mode 0600; these settings prevent others from reading or altering mail while it is in the queue. On the gateway, this has never posed a problem; but it has on the Suns running NFS. Since *root* on a client Sun cannot write to the server's queue directory, Worse, even if it could, the client's *sendmail* process could not read the files it put there! At this point, there are two options: either make the queue files mode 0660 and the queue directory 0777 (meaning anyone can delete queue files, although not necessarily read them), or make each Sun's queue directory a private one (by linking the queue directory to a directory in "/private".) At RIACS, the latter is done; on the server, "/usr/spool/mqueue" is a symbolic link to "/private/usr/spool/mqueue".
- g,u** When *sendmail* spawns a subprocess to deliver mail, it resets the user and group identification numbers to 1 (for *daemon*.) This provides a measure of security; were the mailer running as *root*, a breach of security could result in compromise of the entire system rather than just daemons.
- L** All mail coming into and going out of RIACS is logged; the sender and recipient are named in the log file, as is the queue number and other useful information. To reduce the amount of logging, lower this number. Logging is done via *syslog(3)*, so to determine where information is being sent, look there.
- o** Since many mailers use spaces rather than commas "," to delimit names, this option instructs *sendmail* to accept and handle such lists.
- r** This option instructs *sendmail* to time out after 1 hour during a connection to send mail. Supposedly, this should "never happen", but we all know that the real world rarely lives up to our expectations. With the current state of the internet, a connection is often broken but one or both ends do not know this. Better to time out than wait forever (or for the

next reboot.)

- S** This is where statistics on *sendmail* are kept. The file does not grow, so the “/usr” file system will not overflow due to this.
- s** This is a safety measure; it ensures the queue is always current, even when *sendmail* will deliver the message immediately. With mail, paranoia is a healthy state of mind!
- T** After 3 days, the message will be returned as undeliverable. Another popular value is 15 days.
- t** This is a relic of V6 UNIX and is ignored here.
- W** This option applies to versions of *sendmail* compiled with the “wizard” option (4.3 BSD) or the “debug” option (4.2 BSD.) Those versions of *sendmail* provide a special set of commands, ostensibly for debugging, that allow anyone with access to the SMTP server to break into the system. **DO NOT DELETE THIS OPTION!** With this option set, an attacker cannot break in this way even if the running version of *sendmail* contains the security hole.
- X,x** These control how load averages affect SMTP connections. If the load average is 8 or more, incoming messages are queued and not delivered. If the load average is 12 or more, all requests for an SMTP connection are refused.

### 3.3. Precedences

This section is ignored unless one of the header fields is a “Precedence: ” field, in which case the message is given the appropriate precedence with respect to all other undelivered messages. By default all messages have precedence 0 (the same as “first-class”.) Table 7 summarizes the precedences RIACS recognizes.

Table 7. RIACS Mail Precedence Levels	
class	priority
junk	-100
first-class	0
special-delivery	100

### 3.4. Trusted Users

When forwarding mail, *sendmail* must sometimes specify a different sender than in the message headers (this occurs in intra-network mailing.) These users can use the appropriate command flags to do this; no-one else can. Putting a user in this field means he can change the sender. Currently, the trusted users are *root*, *daemon*, *uucp*, and *network*.

### 3.5. Mailers

This section describes the mailers for the gateway. There are five of them: *local*, which delivers mail with recipients on that host, *prog*, which delivers mail to programs or servers on that host, *uucp*, which delivers mail to the uucp network, *tcp*, which delivers mail over the internet, and *utcp*, which delivers uucp mail over the internet. Non-gateway hosts use three mailers (*local*, *prog*, and *tcp*) with identical descriptions to the mailers for the gateway.

The mailers each have flags indicating what actions *sendmail* should take when invoking them; these all follow the “F=” in each entry. All mailers require “Date: ”, “From: ”, and “Message-Id: ” fields (flags **D**, **F**, and **M**) so if any of those fields are missing *sendmail* will add them before the mailer is invoked. The *local* and *prog* mailers both perform final delivery (the flag **l**); the letter will not be passed to another delivery agent when one of these mailers is invoked. Both the *local* and *tcp* mailers can deliver letters to multiple addresses simultaneously (the flag **m**) so *sendmail* will only issue one command to send the letter, rather than one such command per addressee. UNIX mailers do not conform to the standard, usually requiring headers of the form “From *user* ...” to be the first line in the letter; since the *local* mailer adds those

lines, *sendmail* will not (the flag *n*); also, *sendmail* is to specify the recipient with the *-r* argument to the *local* mailer (the *r* flag). All mailers but the *tcp* mailer require quotes to be stripped from the address before being called, and *sendmail* does this (the *s* flag.) The *prog*, *tcp*, and *utcp* mailers are expensive to connect to, so the gateway will only connect during a queue run; notice the *e* flag. (This flag only has an effect if the "*c*" flag is specified in the Options section. RIACS does not do this.) Also, the *uucp* mailer requires a special line at the top of the letter (this line is of the form "remote from *host* ..." and is also a violation of the standard) so the *U* flag has *sendmail* add it before sending the letter to that mailer. The *uucp*, *tcp*, and *utcp* mailers preserve the case of letters in user names (the *u* flag), and the *uucp* and *utcp* mailers also preserve the case of letters in host names (the *h* flag.) Finally, the *tcp* and *utcp* mailers have a line length limit as specified in RFC 821 [RFC821]; *sendmail* will split lines in the letter if need be to keep lines short enough (the *L* flag.)

Three mailers have other special considerations. Due to limits inherent in the *uucp* network, it is exceedingly unwise to send a letter with more than  $2^{16}-1$  characters; the field "M=65535" in the *uucp* and *utcp* mailers prevents *sendmail* from sending files larger than that through this mailer. The *tcp* and *utcp* mailers use the two-character string "<CR><LF>" (that is, a carriage return followed by a line feed) to signal the end of a line.

The *uucp* mailer is peculiar to 4.3 UNIX. When called, */usr/bin/uuz* is invoked as "*uuz -z -a\$F -gA \$h!rmail (\$u)Q*"; this just means *uuz* will read a letter from the standard input, send error messages to the sender (the *\$F* is replaced with the sender's address), queue the message in the UUCP mail system with high priority (the "*-gA*" does this), and sends the input to the command *rmail* on the remote host *\$h* with the recipient's name in parentheses on the command line.

The *tcp* mailer is really a "pseudo-mailer", because *sendmail* is the delivery agent. The string "*P=[IPC]*" tells *sendmail* to use an SMTP protocol to transmit the message. No other program is invoked (there is no program named "IPC", in fact.) By default, communication is done over port 25; if some other port should be used, name the port after the "*\$h*" in the "*A=*" string; for example, "*A=IPC \$h 100*" initiates contact over port 100.

The *utcp* mailer is a compromise between the *uucp* mailer and the *tcp* mailer. Some sites on the internet are also *uucp* hosts, so rather than queue the message within the *uucp* system and send it that way, we use the SMTP protocol to transmit it directly. Thus, the mailer has the characteristics of both the *uucp* and *tcp* mailers in its description, but uses the *uucp* mailer's rewriting rules to rewrite addresses.

### 3.6. Headers

The header lines indicate what header fields should be added. The "Received:" field is inserted into every message, and the others are inserted depending on what flags the mailers in the *mailers* section above have. Table 5 shows which of the other headers are inserted into letters being handled by the mailers, and the flag present in the "*F=*" field of the mailer description that causes the header line to be inserted.

Table 5. Table of Header Lines		
header line	inserted for these mailers	flag
Return-Path:	none	P
Resent-Date:	local, prog, uucp, tcp	D
Resent-Message-Id:	local, prog, uucp, tcp	M
Message-Id:	local, prog, uucp, tcp	M
Date:	local, prog, uucp, tcp	D
Resent-From:	local, prog, uucp, tcp	F
From:	local, prog, uucp, tcp	F
Full-Name:	none	x

If any of these lines are in the header, another new one will not be added. The test is done on



the flag and not the header itself; for example, if a message has the header field "Date:", neither a new "Date:" nor a "Resent-Date" header field will be added, since a field tied to the D flag is present.

The next two sections describe the rulesets used to rewrite and deliver mail.

#### 4. Sendmail Gateway Configuration File

When *sendmail* obtains an address for processing, it runs that address through a series of rules grouped into rulesets. These rulesets can change the address, pass it on unchanged, or substitute a new address entirely for it. This section describes what RIACS' configuration file does.

##### 4.1. Domain Names and Canonicalization

All the rulesets, except ruleset 3, assume the address is in a standard form; this reduces the number of rules that each ruleset needs. The form assumed is that the domain to which the letter is to be sent is surrounded by angle brackets "<", ">". For example, the *canonical* form of "mab@riacs.edu" is "mab<@riacs.edu>"; the canonical form of "@purdue.edu:mab@riacs.edu" is "<@purdue.edu>:mab@riacs.edu" because the letter is to be sent to "purdue.edu", and from there to "riacs.edu". This form is internal; the angle brackets are removed before the address is written back into the header, or passed to the mailer. It is simply a convenience that allows simpler rewriting rules.

There are some cases where no legitimate domains exist to cover the addresses; for example, UUCP is not a domain in the sense RFC 920 [RFC920] defines domain. But it is much easier to pretend that there is a domain ".uucp", and deal with addresses in that context; we can use the rewriting rules for legitimate domains to process addresses for these sites too. So the address "megatest!mab" is canonicalized to "mab<@megatest.uucp>". Also, many people use UUCP addressing to route letters through Internet sites; such addresses have the form "decwrl.dec.com!megatest!mab". In these cases, the canonicalization process does *not* add a ".uucp" to the domain; it uses the given domain. So, the above address would be canonicalized as "megatest!mab<@decwrl.dec.com>". This approach actually is a bit perilous, because many sites use domain names even though they are not on the Internet; such cases are handled by specific rules affecting only the address handed to a transport mechanism. This way, we can use the same canonical form of an address for all cases.

##### 4.1.1. Ruleset 3

Ruleset 3 canonicalizes all addresses. We shall go through this ruleset in detail, showing how each rule moves the address towards canonicalization. When *sendmail* gets an address, it may be surrounded with "<", ">"; so, we must strip these off. First, we handle the case where no address is provided; then we get the innermost address:

R<>	\$@@
R\$*<\$*<\$*<\$+>\$*>\$*>\$*	\$4
R\$*<\$*<\$+>\$*>\$*	\$3
R\$*<\$+>\$*	\$2

The first rule simply replaces the empty address with an "@" sign, and returns that as the new address. Other rulesets will deal with it later on. The next two rules deal with multiple nestings of addresses; this is not covered by the standard, so it is ambiguous; the convention most places seem to have adopted is to treat the innermost pair of angle brackets as delimiting the address. The final rule simply treats whatever is in angle brackets as being the address.

The next rule,

R\$+ at \$+	\$1@\$2
-------------	---------

simply rewrites an archaic address specification to an acceptable one (RFC 733 [RFC733] allowed it, but RFC 822 [RFC822] does not.) The old form of the address is now illegal, but since many sites will be slow to convert, it still should be recognized.

The next rule converts route specifications into an internal form:

```
R@$, $+           @$1:$2
```

This form replaces all “,”s with “:”s, so (for example) “@site1,@site2:user@site3” would become “@site1:@site2:user@site3”. This is not a legal form, but dealing with this form rather than the legal one makes rulesets much simpler. (The address is rewritten to the legal form before being output.) Note we wish to forward the message to the first host in the specification.

Next, we deal with some of the more obvious errors.

```
R:$+           $$>3$1
R@:$*          $$>3$1
R$*@           $$>3$1
```

These three rules handle the case of a null domain name; in the first case, the separator “:” is deleted, and in the last two cases, the “@” introducing the null domain name is stripped and the address reprocessed.

Now we are ready to canonicalize:

```
R@$:+$+        $$<@$1>:$2
R$+@$+         $:$1<@$2>
R$+<$+@$+>     $1$2<@$3>
R$*<@$*.$*>    $$>3$1@$2$3
R$+<@$+>        $$<@$1>:<@$2>
```

The first rule handles route-specified addresses; note that it returns the result, since later rules shift the angle brackets right and mail is sent to the leftmost host in the route specified addresses rather than the rightmost. The second rule adds angle brackets to other addresses, and the third rule ensures they are placed around the rightmost domain name. The fourth line deals with an error situation; when a domain name ending in a “.” is given, it is an error, so the offending “.” is deleted and the name recanonicalized. Otherwise, the result of the canonicalization is returned.

The last section of this ruleset handles syntaxes of addresses not falling within the scope of [RFC822]. These include UUCP and BITNET (among others.)

```
R$+^$+         $1!$2
R$-!$+         $$2<@$1.uucp>
R$+!$+         $$2<@$1>
R$-!$+         $$2<@$1>
R$-.$+         $$2<@$1>
R$-.$+         $$2<@$1>
R$-=$+         $$2<@$1.bitnet>
R$+%%$+        $$>9$1%$2
```

The first three rules deal with UUCP. The first rule converts an old UUCP address to the new form, and the second and third rules do the conversions described at the beginning of this section. The following two rules convert BERKNET addressing syntax to that of [RFC822]. The rule after that deals with BITNET syntax; like UUCP, BITNET is a pseudo-domain “.bitnet” and is dealt with similarly. The last rule translates “%” to “@” when appropriate; this is done by invoking ruleset 9, which will be discussed next.

#### 4.1.2. Ruleset 9

One very common situation arising on the Internet is that of network routing; for example, if someone at an ARPANET site wishes to send a letter to someone at a CSNET site, he must indicate that the message is to go to another network, the CSNET network. This requires the message to be sent to a site on both ARPANET and CSNET (such sites are called *relay sites* or *relays*.) The obvious syntax is to use routing, as “@relay.cs.net:postmaster@vpi.csnet” but a far more common syntax is to write “postmaster%vpi.csnet@relay.cs.net”. It is more desirable to

have the mail routing mechanism worry about how to get the message from ARPANET to CSNET; with sites that allow this (such as RIACS), people tend to write "postmaster%vpi.csnet", that is, use a "%" rather than a "@". The proper way to deal with such an address is to send the message to the *last* site named. Ruleset 9 deals with addresses involving sequences of "%"; it changes the final such character to "@" and canonicalizes the result:

R\$*%\$*	\$1@\$2
R\$*@\$*@\$*	\$1%\$2@\$3
R\$*@\$*	\$@\$1<@\$2>

The first rule changes all "%" characters in the address to "@" characters; the second rule changes all but the last "@" back. The third rule adds the "<" and ">" around the last site. This canonicalizes the name. Note that if no domain or site name is given after the "%", ruleset 9 is not called, so no error checking need be done. In fact, throughout the other syntaxes ruleset 3 recognizes, there is an implicit assumption that no attempt will be made to recover from a faulty address. Such addresses will cause the mail to fail and be returned to the sender at some point, either at RIACS or at a site farther along the mail path.

## 4.2. The Transport Mechanism

Ruleset 0 is the basis for transport. By the time this ruleset returns, the input address must be resolved to a mailer, a host, and an address. When ruleset 0 returns, *sendmail* instructs the named mailer to contact the named host and send the letter to the named address. In most cases, the mailer is determined by the domain, although there is a mechanism that allows the name of the site to determine the transport mechanism. By default, the pseudo-domain ".uucp" uses the *uucp* mailer; and other domains either use the *tcp* mailer or are rewritten to send to a relay site. The relay site may use either the *tcp* or *uucp* mailers (currently, they all use the *tcp* mailer.) All addresses given to the *tcp* mailer are enclosed in angle brackets and contain the name of the host to which the letter is being given; this, as many hosts reject addresses not in angle brackets or without the name of the host.

### 4.2.1. Ruleset 0

The first few rules of ruleset 0 rewrite the address into a form that can be analyzed and turned into a mailer/host/address set. The first rule eliminates loops back to RIACS; this is an efficiency consideration, and could be dropped.

R\$*<@\$+>\$*	\$\$>6\$1<@\$2>\$3
---------------	--------------------

We shall discuss how this is done when we look at ruleset 6, below. The next rule handles messages with no addresses.

R@	\$#local\$:\$n
----	----------------

Recall that ruleset 3 changed null addresses into "@"; this rule just sends the letter to the designated person: The next rule deals with domains to which RIACS does not have direct access:

R\$*<@\$+>\$*	\$\$>8\$1<@\$2>\$3
---------------	--------------------

These must be reached through relays; all this resolution is done in ruleset 8, which is used to rewrite addresses here. Note that some subdomains of known domains are rewritten; this need not be done once RIACS uses a name resolver, but should be kept just in case the name resolver is replaced. The rewriting must be done if the top-level domain is one RIACS can only access through a relay (such as BITNET or CSNET.)

Finally we must convert route specification addresses from the internal form to the legal form. This rule does so:

R\$*:@\$*	\$1,@\$2
-----------	----------

Wherever the sequence ":@\*" appears, it is replaced by "@"; this converts the internal format of

a route specified address to the legal form, which can then be output.

The address is now in a form that can be resolved to a mailer/host/address triple. Some RIACS hosts do not accept SMTP connections, and others are not accessible via *uucp*; table 6 summarizes the mailers used to reach each local host. They are divided into two classes based on the mailer used.

Table 6. Mailers for RIACS Local Hosts					
host	mailer	class	host	mailer	class
clavier	tcp	T	icarus	tcp	T
cube	uucp	U	lavalite	tcp	T
daedalus	tcp	T	miranda	tcp	T
dora	tcp	T	phun	tcp	T
hydra	tcp	T	zeus	tcp	T

Rather than being rigid and rejecting all incorrectly-routed mail, mail is rerouted properly. The next four rules do this:

```

R$*<@=T.uucp>$*      $#tcp$@2$:<$1@2$3>
R$*<@=U>$*            $#uucp$@2$:$1$3
R$*<@=U.N>$*          $#uucp$@2$:$1$3
R$*<@=U.arpa>$*       $#uucp$@2$:$1$3

```

The first rule changes any mail sent to a RIACS SMTP host over *uucp* to use the SMTP mailer *tcp*. The next three rules change any mail sent to a RIACS host that does not accept mail from the *tcp* mailer to use *uucp*.

One last problem remains. Many sites in the "uucp" (and other) domains use the ARPA domains, and so mail must be routed to them on a per-site basis. Also, many "uucp" sites use the ARPANET as the communications medium for *uucp*; it makes more sense to send mail over the ARPANET rather than use *uucp*. This ruleset is used to do such routing. No mailers are invoked; the addresses are simply rewritten. Table 7 describes the sites the addresses of which are rewritten:

Table 7. Addresses that RIACS Rewrites			
site names		mailers	
incoming	outgoing	original	replaced by
amdcad.amd.com	amdcad	tcp	uucp
ames	ames.arc.nasa.gov	uucp	utcp
decwrl	decwrl.dec.com	uucp	utcp
ll-xn	ll-xn.arpa	uucp	utcp
lll-crg	lll-crg.arpa	uucp	utcp
rutgers	rutgers.rutgers.edu	uucp	utcp
seismo	seismo.css.gov	uucp	utcp

The rules are:

```

R$*<@amdcad.amd.com>$*      $#uucp$@amdcad$:$1$2
R$*<@ames.uucp>$*           $#utcp$@ames.arc.nasa.gov$:<$1@ames.arc.nasa.gov$2>
R$*<@decwrl.uucp>$*         $#utcp$@decwrl.dec.com$:<$1@decwrl.dec.com$2>
R$*<@ll-xn.uucp>$*          $#utcp$@ll-xn$:$1@ll-xn$2>
R$*<@lll-crg.uucp>$*        $#utcp$@lll-crg$:$1@lll-crg$2>
R$*<@rutgers.uucp>$*        $#utcp$@rutgers.rutgers.edu$:<$1@rutgers.rutgers.edu$2>
R$*<@seismo.uucp>$*         $#utcp$@seismo.css.gov$:<$1@seismo.css.gov$2>

```

More rules can be added as needed.

We next dispose of addresses within the "uucp" pseudo-domain:

```
R<@$.uucp>:$+      $#uucp$@1$:$2
R$+<@$.uucp>      $#uucp$@2$:$1
```

Note that these handle "uucp" addresses in both route specifications and other forms. Two rules are necessary because the name of the machine to which the letter is being sent is not included in the address; for example, a letter to "megatest!ametek!root" would call the *uucp* mailer with a host of "megatest" and an address of "ametek!root".

The next rule handles all addresses for the *tcp* mailer:

```
R$*<@$*>$*      $#tcp$@2$:<$1@2$3>
```

Note that the address handed to the mailer is enclosed in angle brackets. This handles both rout-specified addresses and other forms of addressing. Finally, any address not converted to a mailer/host/address triple is intended to be delivered on this machine, so the final rule in ruleset 0 does this:

```
R$+      $#local$:$1
```

#### 4.2.2. Ruleset 6

Earlier we mentioned some other rulesets; let us now look at ruleset 6. This ruleset eliminates the gateway as a target for a letter; otherwise, the gateway would just deliver mail to itself, a rather pointless exercise especially when it can be handled immediately. The rules in this ruleset are actually divided into two groups: the ones that do canonicalization, and the ones that recurse. The first six rules all recanonicalize after stripping the gateway's name; the next six do not change the address, but reinvoke this ruleset if the target host is the gateway.

The first two rules delete the top-level domain and uucp names from the address:

```
R$*<@$N>$*      $>3$1$2
R$*<@$U>$*      $>3$1$2
```

Note these simply delete the host name in angle brackets and recanonicalize. The next four deal with the name of the gateway:

```
R$*<@$=w>$*      $>3$1$3
R$*<@$=G>$*      $>3$1$3
R$*<@$=G.$N>$*    $>3$1$3
R$*<@$=G.arpa>$*  $>3$1$3
```

Because the gateway is known by so many different names, and the domain may or may not be appended, the first two rules deal with the gateway host name without the domain, the third with the fully qualified gateway host name, and the last with the old ".arpa" form of the gateway host name. The next six rules are similar, but just reinvoke this ruleset:

```
R$*<@$N>$*      $>3$1<@$N>$2
R$*<@$U>$*      $>3$1<@$U>$2
R$*<@$=w>$*      $>3$1<@$2>$3
R$*<@$=G>$*      $>3$1<@$2>$3
R$*<@$=G.$N>$*    $>3$1<@$2.$N>$3
R$*<@$=G.arpa>$*  $>3$1<@$2.arpa>$3
```

They are present in case the recanonicalization has focused on another version of the gateway's name.

#### 4.2.3. Ruleset 8

Ruleset 8 deals with the domains to which RIACS does not have direct access. This requires the address to be rewritten in order to send the message through a relay host. There are two syntaxes used; the first, for unrouted addresses, is to replace the address with something the relay can use. For example, the address "root@munniari.oz" would become

"root%munari.oz@seismo.css.gov" since "seismo.css.gov" is the relay host for the domain "oz". The relevant rules are:

R\$*<@decwrl.dec.com>\$*	\$Q\$1<@decwrl.dec.com>\$2
R\$*<@decwrl.dec>\$*	\$Q\$1<@decwrl.dec.com>\$2
R\$*<@\$+.au>	\$Q\$1%\$2.au<@seismo.css.gov>
R\$*<@\$+.bitnet>	\$Q\$1%\$2.bitnet<@wiscvm.wisc.edu>
R\$*<@\$+.csnet>	\$Q\$1%\$2.csnet<@relay.cs.net>
R\$*<@\$+.dec.com>	\$Q\$1%\$2.dec<@decwrl.dec.com>
R\$*<@\$+.dec>	\$Q\$1%\$2.dec<@decwrl.dec.com>
R\$*<@\$+.decnet>	\$Q\$1%\$2.decnet<@\$ ames-io\$ >
R\$*<@\$+.mailnet>	\$Q\$1%\$2.mailnet<@\$ mit-multics\$ >
R\$*<@\$+.oz>	\$Q\$1%\$2.oz<@seismo.css.gov>
R\$*<@telemail>	\$Q\$1%telemail<@orion.arc.nasa.gov>

The first two rules make sure that letters being sent to "*someone*@decwrl.dec.com", the DEC Easynet relay, are not addressed to "*someone*%decwrl.dec@decwrl.dec.com" (for some reason, "decwrl.dec.com" cannot handle this.) The remaining nine rules simply rewrite the addresses as required. Be aware that the rules for "dec.com" will go away when the name resolver software becomes reliable; in fact, since RIACS has the table of DEC hosts, all rules involving "dec" (except which is an Ames local network) may soon be commented out.

The second syntax is used for route-specified addresses; the relevant rules are:

R<@\$+.au>\$*	\$Q<@seismo.css.gov>:\$1.au:\$2
R<@\$+.bitnet>\$*	\$Q<@wiscvm.wisc.edu>:@\$1.bitnet:\$2
R<@\$+.csnet>\$*	\$Q<@relay.cs.net>:@\$1.csnet:\$2
R<@\$+.dec.com>\$*	\$Q<@decwrl.dec.com>:@\$1.dec:\$2
R<@\$+.dec>\$*	\$Q<@decwrl.dec.com>:@\$1.dec:\$2
R<@\$+.decnet>\$*	\$Q<@\$ ames-io\$ >:@\$1.decnet:\$2
R<@\$+.mailnet>\$*	\$Q<@\$ mit-multics\$ >:@\$1.mailnet:\$2
R<@\$+.oz>\$*	\$Q<@seismo.css.gov>:\$1.oz:\$2
R<@telemail>\$*	\$Q<@orion.arc.nasa.gov>:@telemail:\$2

The rules just put the relay before the host to which the letter is to be directed; for example, the address "<@munari.oz>:someone@otherhost" will be rewritten as "<@seismo.css.gov>:@munari.oz:someone@otherhost". Notice that addresses to "decwrl.dec.com" are taken care of by the first two rules in the ruleset, and so need not be repeated here.

#### 4.3. The Sender's Address

There are two philosophies for handling a sender's address. Either the address is rewritten to reflect the way the message is sent, or it is not rewritten unless failing to do so would produce known errors when that path is used as a reply path. RIACS takes the latter approach. We assume that other sites know best how to route their mail, and so (with the exception of *uucp*, for a reason explained below) we do not rewrite the sender's address unless the sender is at this site.

The rewriting rulesets applied to the sender's address depends in part on the mailers being used. In all cases, ruleset 1 is applied; then the mailer-dependent rules; and finally ruleset 4 operates on the result.

##### 4.3.1. Ruleset 1

Ruleset 1 hides the name of the local host within RIACS by replacing it with the name of the gateway as a domain; so, "root@hydra" would be rewritten as "root@riacs.edu". This enables RIACS to hide its internal configuration from the rest of the world; it is most useful should the internal configuration change, since only the gateway would need to be updated.

Table 8 summarizes this configuration; note the nicknames do not have any domain or pseudo-domain names appended.

Table 8. RIACS Local Hosts	
official host name	nicknames
clavier.riacs.edu	clavier riacs-clavier
cube.riacs.edu	cube hypercube riacs-cube riacs-hypercube
daedalus.riacs.edu	daedalus riacs-daedalus ames-daedalus
dora.riacs.edu	dora riacs-dora
hydra.riacs.edu	hydra riacs-hydra
icarus.riacs.edu	icarus riacs riacs-icarus riacs-gw
lavalite.riacs.edu	lavalite riacs-lavalite
miranda.riacs.edu	miranda riacs-miranda
pegasus.riacs.edu	pegasus riacs-pegasus ames-pegasus
phun.riacs.edu	phun riacs-phun
zeus.riacs.edu	zeus riacs-zeus

This ruleset is quite simple, consisting of only four rules:

```

R$*<@H>$*           $$1<@N>$3
R$*<@$=H.arpa>$*     $$1<@N>$3
R$*<@$=H.N>$*        $$1<@N>$3
R$*<@$=H.uucp>$*     $$1<@U>$3

```

The first rule changes unqualified names, the second changes old-style ".arpa" names, and the third changes fully qualified hostnames. Notice the fourth, which translates local host names into the standard RIACS *uucp* designator.

#### 4.3.2. Ruleset 10

The *local* and *prog* mailers next use ruleset 10, which consists of one rule:

```

R@                $n

```

If there is a null address, this ensures the appropriate person gets the (rejected) letter; that person can then decide what action to take.

#### 4.3.3. Ruleset 13

The *uucp* mailer user ruleset 13. This ruleset must prepend the RIACS *uucp* name *riacs.uucp* to the address, so (for example) "mab@<megatest.uucp>" would become "megatest!mab<@riacs.uucp>"; if this is not done, the remote *uucp* mail handler will record the return address as "somewhere!mab" rather than "riacs!mab", making it impossible to reply to that letter. Five rules accomplish this:

```

R$*<@U>$*           $$1<@U>$2
R$*<@$.uucp>        $$2!$1<@U>
R<@$.uucp>$*        $$<@U>:@$1.uucp:$2
R$*<@$*>$*          $$1<@2>$3
R$*                  $$1<@U>

```

The first rule simply eliminates cases where *riacs.uucp* is already in the address. The next two rulesets add the name to addresses sent via *uucp* and via routing specifications. The fourth rule returns any nonlocal addresses; presumably, these are of the form "*someone@amdcad.amd.com*" where RIACS can only talk to "amdcad.amd.com" via *uucp*. In this case, the remote site's *uucp* mail agent must be able to handle the address; RIACS should not have to. The final rule adds the *uucp* domain to addresses from a local machine; so, "mab" would be rewritten as "mab<@riacs.uucp>".

#### 4.3.4. Ruleset 14

The *tcp* mailer uses ruleset 14. If the sender is local, this ruleset appends the RIACS domain to the address:

R\$*<@\$*>\$*	\$Q\$1<@2>\$3
R\$+	\$\$1<@N>

#### 4.3.5. Ruleset 4

Finally, ruleset 4 does some miscellaneous cleanup:

R@	\$@
----	-----

Just in case an empty address gets this far, it is ignored; the next mailer can deal with it. (This should never happen, but just in case this line is left here.) The next rule decanonicalizes the address by deleting the angle brackets "<", ">":

R\$*<\$+>\$*	\$1\$2\$3
--------------	-----------

Route specified addresses must be completely surrounded by such brackets, and converted to a legal form; the next rules do this:

R\$+:@\$+	\$1,@2
RQ\$+:\$+@\$+	<@1:\$2@3>

and the final rule rewrites uucp addresses to their usual form:

R\$*@\$+.uucp	\$2!\$1
---------------	---------

Note this will not affect uucp addresses which are part of a route specification.

### 4.4. The Recipient's Address

As with senders' addresses, there are two philosophies for handling a recipient's address. Either the recipient's address is not rewritten unless failing to do so would produce known errors when that path is used by the next host to which the message is sent, or the address is rewritten to reflect the routing used to get to the next host in the address. RIACS takes the former approach. We assume that other sites know best how to route their mail, and so we do not rewrite the recipient's address unless the recipient is at this site.

The rewriting rulesets applied to the recipient's address depends in part on the mailers being used. In all cases, ruleset 2 is applied; then the mailer-dependent rules; and finally ruleset 4 operates on the result.

#### 4.4.1. Ruleset 2

*Sendmail* runs the address given to the mailer through these rulesets, so ruleset 2 is used to recanonicalize these addresses. Note this also works for recipient addresses handed directly to ruleset 2 (that is, those from the "To:" or "cc:" field) because defocusing and recanonicalizing is essentially a no-op in this case. Ruleset two has two rules:

R\$*<\$*@\$*>\$*	\$\$1\$2\$3\$3
R\$*	\$Q\$>3\$1

The first rule deletes all angle brackets, and the second rule just invokes ruleset 3.

#### 4.4.2. Ruleset 20

The *local* and *prog* mailers next use ruleset 20, which needs to do nothing; so there are no rules in it.



#### 4.4.3. Ruleset 23

The *uucp* mailer uses ruleset 23, which (like ruleset 20) needs to do nothing; so there are no rules in it.

#### 4.4.4. Ruleset 24

The *tcp* mailer uses ruleset 24. If the recipient is local, this ruleset appends the RIACS domain to the address:

R\$*<@\$*>\$*	\$\$>8\$1<@\$2>\$3
R\$+	\$\$1<@\$N>

Ruleset 4 has been discussed above.

#### 4.5. Conclusion

This completes the discussion of why the gateway configuration file was set up as it is. Now let us look at the non-gateway file.

### 5. Sendmail Non-Gateway Configuration File

This file is much simpler, because the strategy is simply to send any nonlocal mail to the gateway. Accordingly, sender and recipient addresses are not processed, and canonicalization is done only to aid in determining the mailer to be used.

#### 5.1. The Transport Mechanism

Rulesets 3 and 9 are the same as for the gateway configuration file, so they will not be discussed again; and ruleset 0 is much simpler.

##### 5.1.1. Ruleset 0

As with the gateway, the first rule in that ruleset ensures all names of the host are deleted by invoking ruleset 6:

R\$*	\$\$>6\$1
------	-----------

At this point, the mailer can be determined. Either the letter is local, in which case the *local* or *prog* mailers need to be invoked, or it is not local, in which case the *tcp* mailer sends the letter to the gateway:

R@	\$\$local\$:\$n
R\$*<\$*@\$*>\$*	\$\$tcp\$@\$G\$:<\$1\$2@\$3\$4>
R\$+	\$\$local\$:\$1

Note that if the address is empty, the first rule will send the letter to the appropriate user, who can take whatever action is deemed appropriate. The next rule just forwards any mail with a host name to the gateway; since names for this host were deleted by the first rule in this ruleset, such mail is destined for another host. The last rule just delivers the mail locally.

##### 5.1.2. Ruleset 6

Ruleset 6 is somewhat different because it does not have to deal with domain names or "uucp" names, but only the names of this machine. As in the gateway configuration file, the first set of rules recanonicalize the address, and the next set determines whether or not to reinvoke ruleset 6:

R\$*<@\$=w>\$*	\$\$>3\$1\$3
R\$*<@\$=w.arpa>\$*	\$\$>3\$1\$3
R\$*<@\$=w.\$N>\$*	\$\$>3\$1\$3
R\$*<@\$=w>\$*	\$\$>6\$1<@\$2>\$3
R\$*<@\$=w.arpa>\$*	\$\$>6\$1<@\$2.arpa>\$3

R\$\*<@\$.w.\$N>\$\*

\$>6\$1<@\$.2.\$N>\$3

## 5.2. Philosophy

The philosophy in writing the non-gateway configuration files has been to keep the rules as simple as possible. This is why local hosts will send mail to each other through the gateway rather than directly; since not all hosts can speak SMTP, only the gateway needs to keep track of how each host is to receive mail. This makes the gateway file the only one that needs to be changed when the configuration changes.

## 6. Debugging a *Sendmail* Configuration File

RIACS runs various versions of *sendmail* on its hosts. All are essentially the same program, so the techniques discussed below work for all versions of *sendmail*. In all cases, it is **strongly** recommended that the new configuration file **not** be installed until it is fully tested and debugged; instead, run *sendmail* with the **-C** option. So, to debug a configuration file *new.cf*, issue the command

`sendmail -Cnew.cf other options ...`

The first mode for testing is called *test mode*, appropriately enough, and is used to run addresses through rulesets to show what each ruleset is given, what it returns, and how the given address is rewritten. To use it, put the flag **-bt** on the command line. Here is a sample session, using the gateway configuration file; some of the longer lines have been split into two, in order to fit on the page.

```
% /usr/lib/sendmail -bt
ADDRESS TEST MODE
Enter <ruleset> <address>
> 0 riacs.edu!site.com!xyz
rewrite: ruleset 3 input: "riacs" "." "edu" "!" "site" "." "csnet" "!" "xyz"
rewrite: ruleset 3 returns: "site" "." "csnet" "!" "xyz" "<" "@ "riacs" "."
                        "edu" ">"
rewrite: ruleset 0 input: "site" "." "csnet" "!" "xyz" "<" "@ "riacs" "."
                        "edu" ">"
rewrite: ruleset 6 input: "site" "." "csnet" "!" "xyz" "<" "@ "riacs" "."
                        "edu" ">"
rewrite: ruleset 3 input: "site" "." "csnet" "!" "xyz"
rewrite: ruleset 3 returns: "xyz" "<" "@ "site" "." "csnet" ">"
rewrite: ruleset 6 returns: "xyz" "<" "@ "site" "." "csnet" ">"
rewrite: ruleset 8 input: "xyz" "<" "@ "site" "." "csnet" ">"
rewrite: ruleset 8 returns: "xyz" "%" "site" "." "csnet" "<" "@ "relay" "."
                        "cs" "." "net" ">"
rewrite: ruleset 0 returns: "$#" "tcp" "$@" "relay" "." "cs" "." "net" "$:"
                        "<" "xyz" "%" "site" "." "csnet" "@ "relay" "."
                        "cs" "." "net" ">"
> 1,14,4 cde@phun.riacs.edu
rewrite: ruleset 3 input: "cde" "@ "phun" "." "riacs" "." "edu"
rewrite: ruleset 3 returns: "cde" "<" "@ "phun" "." "riacs" "." "edu" ">"
rewrite: ruleset 1 input: "cde" "<" "@ "phun" "." "riacs" "." "edu" ">"
rewrite: ruleset 1 returns: "cde" "<" "@ "riacs" "." "edu" ">"
rewrite: ruleset 14 input: "cde" "<" "@ "riacs" "." "edu" ">"
rewrite: ruleset 8 input: "cde" "<" "@ "riacs" "." "edu" ">"
rewrite: ruleset 8 returns: "cde" "<" "@ "riacs" "." "edu" ">"
rewrite: ruleset 14 returns: "cde" "<" "@ "riacs" "." "edu" ">"
rewrite: ruleset 4 input: "cde" "<" "@ "riacs" "." "edu" ">"
rewrite: ruleset 4 returns: "cde" "@ "riacs" "." "edu"
> 2,24,4 ghi@upi.csnet
rewrite: ruleset 3 input: "ghi" "@ "site" "." "au"
rewrite: ruleset 3 returns: "ghi" "<" "@ "site" "." "au" ">"
rewrite: ruleset 2 input: "ghi" "<" "@ "site" "." "au" ">"
rewrite: ruleset 3 input: "ghi" "@ "site" "." "au"
```

```

rewrite: ruleset 3 returns: "ghi" "<" "@ "site" "." "au" ">"
rewrite: ruleset 2 returns: "ghi" "<" "@ "site" "." "au" ">"
rewrite: ruleset 24 input: "ghi" "<" "@ "site" "." "au" ">"
rewrite: ruleset 8 input: "ghi" "<" "@ "site" "." "au" ">"
rewrite: ruleset 8 returns: "ghi" "% "site" "." "au" "<" "@ "seismo" "."
                           "css" "." "gov" ">"
rewrite: ruleset 24 returns: "ghi" "% "site" "." "au" "<" "@ "seismo" "."
                           "css" "." "gov" ">"
rewrite: ruleset 4 input: "ghi" "% "site" "." "au" "<" "@ "seismo" "."
                           "css" "." "gov" ">"
rewrite: ruleset 4 returns: "ghi" "% "site" "." "au" "@ "seismo" "." "css"
                           "." "gov"

```

In this example, the first ruleset called is ruleset 3 (this is always the case, regardless of what rules you name). It is passed the line "riacs.edu!site.csnet!xyz", and returns the canonicalized form "site.csnet!xyz<@riacs.uucp>". Ruleset 0 is then invoked; its first rule calls ruleset 6, which strips off the "@riacs.edu" (since this example was run on *icarus*, which is also known as *riacs.edu*, there is no point whatsoever in remailing the message to *riacs.edu*) and calls ruleset 3 on the remainder of the string, "site.csnet!xyz". Ruleset 3 rewrites this as "xyz<@site.csnet>" and returns to ruleset 6, which returns control to ruleset 0. Ruleset 0 then calls ruleset 8, which rewrites the address to send it through the appropriate relay ("xyz%site.csnet<@relay.cs.net>") and then returns this new form to ruleset 0. Ruleset 0 rewrites the address as "\$#tcp\$@relay.cs.net\$:<xyz%site.csnet@relay.cs.net>" and using this form calls the appropriate mail handling agent. The other two examples work in the same way, the second example using rulesets 3, 1, 14, and 4, and the third example using rulesets 3, 2, 24, and 4.

The version of *sendmail* on *icarus*, the RIACS gateway, has been modified in two ways. The first, relevant to debugging, prints control strings used in the configuration file as they are entered in that file [GILL86]. This modification has not been made to any other version of *sendmail*, so beware! Usually, the final line of the first case in the above example would be printed as:

```

rewrite: ruleset 0 returns: "^V" "tcp" "^W" "site" "." "com" "^X" "<" "xyz" "@ "
                           "site" "." "com" ">"

```

for 4.3 BSD's version of *sendmail*, or

```

rewrite: ruleset 0 returns: "^U" "tcp" "^V" "site" "." "com" "^W" "<" "xyz" "@ "
                           "site" "." "com" ">"

```

for other versions of *sendmail*. Table 9 shows the internal symbols corresponding to the special symbols used in rewriting rules; the version of *sendmail* on *icarus* uses those in the 4.3 BSD column, and all other machines use those in the 4.2 BSD column.

Table 9. Table of Configuration File Actions								
Symbol	Shown		Symbol	Shown		Symbol	Shown	
	4.2BSD	4.3BSD		4.2BSD	4.3BSD		4.2BSD	4.3BSD
\$*	^P	^P	\$	^T	^U	\$?	^Y(?)	^Z
\$+	^Q	^Q	\$#	^U	^V	\$	^Z(?)	^
\$-	^R	^R	\$@	^V	^W	\$.	^[(?)	^\ none
\$=	^S	^S	\$:	^W	^X	\$[	none	^
\$~	^\ ^T	^T	\$>	^X	^Y	\$]	none	^^

There are also a set of debugging flags that are quite useful when one particular aspect of *sendmail* needs to be examined very closely. They differ from version to version; the ones listed in Table 10 are for version 5.51.

Table 10. Table of Debugging Flags

<i>flag</i>	<i>level</i>	<i>information about</i>
0	1	main: run as daemon in foreground
0	4	main: show canonical name and aliases of current host
0	15	main: print configuration table as loaded from configuration file
0	44	main: print command-line arguments to program
1	1	main: report sender of letter
1	9	main: allow user doing debugging to substitute new sender
2	1	exit: report exit status, flags
5	4	event: report alarms
5	5	event: report setting, clearing of event
5	6	event: show event processing
6	1	error: report how mail is dealt with on error
6	5	error: display state of <i>sendmail</i> on error
7	1	queueing: show names of queue file
7	2	queueing: show name of temporary queue file ("tf" file)
7	20	queueing: report processing of the queued file
10	1	deliver: show address given to mailer
11	1	deliver: show address to which letter is actually sent
12	1	parseaddr: show resolution of remote name in address
13	1	deliver: show to whom messages are to be sent
13	3	deliver: check for errors in sending messages
13	4	deliver: report to whom errors go
14	2	header: print list of names in header with commas
15	1	daemon (server): report requests
15	2	daemon (server): report forking to process a request
15	15	daemon (server): put network in debugging mode
16	1	daemon (client): report making remote connection
16	14	daemon (client): put network in debugging mode
18	1	smtp: report result of trying to make connection
18	100	smtp: pause after error in reading from remote connection
20	1	parseaddr: report address to be parsed
21	2	parseaddr: report input, output of each ruleset
21	3	parseaddr: report call to another ruleset (via "\$>")
21	4	parseaddr: report result of call to another ruleset
21	10	parseaddr: report failure of ruleset to match
21	12	parseaddr: report ruleset to be matched and if match succeeds
21	15	parseaddr: report substitution due to matches substitution
21	35	parseaddr: show attempts to match
22	36	parseaddr: show address prescan
22	45	parseaddr: show what is being prescanned
22	101	parseaddr: show states during parsing of address
25	1	recipient: show recipient list (multiple addresses possible)
26	1	recipient: show an individual address
27	1	alias: report alias definition
30	1	smtp: report end of headers
30	2	smtp: report addition of "Apparently-To:" field
30	3	smtp: report processing of UNIX "From" line
31	6	header: show header line being read
32	1	header: show all headers to current letter
33	1	header: show header line as processed
35	9	macro: report definition
35	24	macro: show macro expansion
36	5	symbols: report addition and/or lookup

Table 10. Table of Debugging Flags

flag	level	information about
36	9	symbols: display symbol hash
37	1	configuration: show option settings in configuration file
40	1	queue: show queue run of file
40	4	queue: print queue file lines ("qf" file) as read
41	2	queue: report problem reading queue during ordering
45	1	envelope: report sender
50	1	envelope: report deallocation
51	4	queue: show deletion of transcript file
52	1	main: report disconnection from controlling terminal
52	5	main: do not disconnect from controlling terminal

Options are set with the `-dflag`, which has the format `-dflaglist.level`. To set flags 5, 6, 7, and 36 at level 10, for example, give the option `"-d5-7,36.10"`. It is usually advisable to confine debugging flags to number 21 (which deals with the way addresses are parsed) or display all of them; also, specifying a level prints all messages from lower levels, too. To print any information that could be of importance, use `"-d0-99.99"`; the levels higher than 99 are used to debug *sendmail's* internals.

As stated above, debugging flag 21 is useful enough to merit some special discussion. When more detail about the address parsing is needed, this flag prints each step in the ruleset and indicates how a match is being made. For example,

```
% /usr/lib/sendmail -bt -d21.99
Version 5.51
ADDRESS TEST MODE
Enter <ruleset> <address>
> 0 mab@riacs.edu
rewrite: ruleset 3 input: "mab" "@" "riacs" "." "edu"
-----trying rule: "<" ">"
ap="mab", rp="<"
----- rule fails
-----trying rule: "$+" "<" "$+" "<" "$+" "<" "$+" ">" "$+" ">" "$+" ">" "$+"
ap="mab", rp="$+"
ap="mab", rp="<"
ap="@", rp="<"
ap="riacs", rp="<"
ap=".", rp="<"
ap="edu", rp="<"
ap=<null>, rp="<"
----- rule fails
-----trying rule: "$+" "@" "$+"
ap="mab", rp="$+"
ap="@", rp="@"
ap="riacs", rp="$+"
ap=".", rp=<null>
ap="edu", rp=<null>
-----rule matches: "$:" "$1" "<" "@" "$2" ">"
$1: 7ffffe2bc="mab"
$2: 7ffffe2c2="riacs" 7ffffe2c8="." 7ffffe2ca="edu"
rewritten as: "mab" "<" "@" "riacs" "." "edu" ">"
rewrite: ruleset 3 returns: "mab" "<" "@" "riacs" "." "edu" ">"
rewrite: ruleset 0 input: "mab" "<" "@" "riacs" "." "edu" ">"
-----trying rule: "$+" "<" "@" "$+" ">"
ap="mab", rp="$+"
ap="mab", rp="<"
ap="<", rp="<"
ap="@", rp="@"
ap="riacs", rp="$+"
ap=".", rp=">"
```

```
ap="edu", rp=">"
ap=">", rp=">"
-----rule matches: "$:" "$>" "6" "$1" "<" "@@" "$2" ">"
$1: 7fffe2bc="mab"
$2: 7fffe2c2="riacs" 7fffe2c8="." 7fffe2ca="edu"
-----callsubr 6
rewrite: ruleset 6 input: "mab" "<" "@@" "riacs" "." "edu" ">"
-----trying rule: "$+" "<" "@@" "riacs" "." "edu" ">"
ap="mab", rp="$+"
ap="<", rp="<"
ap="@" , rp="@"
ap="riacs", rp="riacs"
ap=".", rp="."
ap="edu", rp="edu"
ap=">", rp=">"
-----rule matches: "$>" "3" "$1"
$1: 7fffe2bc="mab"
-----callsubr 3
rewrite: ruleset 3 input: "mab"
rewrite: ruleset 3 returns: "mab"
rewritten as: "mab"
rewrite: ruleset 6 returns: "mab"
rewritten as: "mab"
-----trying rule: "$+"
ap="mab", rp="$+"
-----rule matches: "$#" "local" "$:" "$1"
$1: 7fffe2bc="mab"
rewritten as: "$#" "local" "$:" "mab"
rewrite: ruleset 0 returns: "$#" "local" "$:" "mab"
>
```

(Note that parts of the above were edited to keep the output brief.) When ruleset 3 is first called, it tests to see if the address matches "<>". It does not, as the token "mab" is not the same as the token "<". So it tries the next rule in the ruleset, and continues, until it finds a match in the rule "\$+@\$+". This rule replaces the previous form with "mab<@riacs.edu>", and continues the ruleset; finally, it returns with "mab<@riacs.edu>". Ruleset 0 then calls ruleset 6, which strips off the "<@riacs.edu>" and calls ruleset 3 to recanonicalize the rest of the address "mab". Eventually, ruleset 6 returns "mab", and ruleset 0 turns this into the appropriate mailer command.

Occasionally a user will report a problem with mail requiring the maintainer of the mail system to determine at which hosts, if any, the problem occurred. To determine this, look at the headers in the letter that posed the problem. If there are error messages, read them first; if those do not provide enough information, scan the "Received: " lines. Lines added at RIACS indicate from whom the mail routers think the message was received; this can be used to track backwards. As an example, the message with the headers

```
Received: from icarus.riacs.edu (icarus.ARPA) by hydra.riacs.edu (4.12/1.6N)
        id AA14571; Fri, 16 Jan 87 16:41:02 pst
Received: from RELAY.CS.NET by icarus.riacs.edu (5.51/1.6G)
        id AA00679; Fri, 16 Jan 87 16:40:49 PST
Received: from icarus.riacs.edu by RELAY.CS.NET id aa02271; 16 Jan 87 19:30 EST
Received: by icarus.riacs.edu (5.51/1.6G)
        id AA00630; Fri, 16 Jan 87 16:29:16 PST
Message-Id: <8701170029.AA00630@icarus.riacs.edu>
Date: Fri, 16 Jan 87 16:29:16 PST
From: Matt Bishop <mab@riacs.edu>
To: mab%riacs.edu@RELAY.CS.NET
Subject: show how "Received: " header lines work
```

went from *icarus* (*icarus.riacs.edu*) to *RELAY.CS.NET* and back to *icarus* and from there to *hydra*.

Incidentally, the version of *sendmail* on *icarus* has been modified to report the name of the system from which *icarus* receives the letter. This modification defines the unsupported macro *s* so that the name of the sending host may be included on the "Received" header line, as well as allowing the receiving host to recognize nicknames as well as official names (these two fixes are from [LOVS86].) The *sendmail* files which have been altered are stored in their original form as *file.orig* in the *sendmail* source directory.

## 7. Installing New *Sendmail* Configuration Files

In general, to install a new *sendmail* configuration file, do the following:

1. Locate and kill any currently-running invocations of *sendmail*. This is best done by typing "ps aux | grep sendmail" as *root* and then killing all the processes this command lists. If some have died anyway, do not worry; they exited between the "ps" and the "kill".
2. Copy "/usr/lib/sendmail.cf" somewhere. This way, if the new configuration file does not work right, the old one can be put back.
3. Move the new configuration file to "/usr/lib/sendmail.cf".
4. Freeze this file. This puts it into a form *sendmail* can load quickly. Type "/usr/lib/sendmail -bz"; the file that is created is "/usr/lib/sendmail.fc".
5. Restart the *sendmail* daemon. This command will be of the form "/usr/lib/sendmail -bd -q1h", but it varies from system to system. Look in "/etc/rc.local" for the exact command; typing "grep 'sendmail.\*-bd' /etc/rc.local" will print the exact command.
6. Type "mailq". On the server Sun, the command will print a message saying "Freeze file out of date", and delete the frozen configuration file; this is normal. (Why it is done is not known.) This ensures the *sendmail* program is working correctly.

On the client Suns, the procedure is slightly different since they use the server's configuration file. On these systems, omit steps 2 through 4.

## 8. Changing the Configuration Files

This section describes how to make routine changes to the *sendmail* configuration files. It does not cover massive rewriting; for that, your best bet is to find a *sendmail* guru or just experiment.

### 8.1. Adding a New RIACS Host (Gateway Configuration File)

This is really quite easy. First, figure out how mail will be sent to the host; if it will go via *uucp* or *tcp*, you're in luck. (If not, you'll have to define a new mailer. See [ALLM84b] for a description of how to do this; use the already-created mailer descriptions as a guide.)

First, add all names of the new host in the class **H**; so if the new host were named "hera.riacs.edu", with nicknames "riacs-hera.arpa", "hera", "ames-hera", and "riacs-hera", add a line of the form

```
CHhera riacs-hera ames-hera
```

to the appropriate section (where the rest of the hosts are defined.) Then, if the mailer is *tcp*, add all the names to the class **T**; if the mailer is *uucp*, add all the names to the class **U**. If you are using some other mailer, define a new class for the mailer (remember, *sendmail* macro names are one character only!), and put this host into the class. For explanatory purposes, call the new class **W** and the mailer *newmail*. Then in ruleset 0, add the following rules to the end of the "resolve the mailers ... handle local hosts" section:

```
R$*<$=W>$*          $#newmail$@2$:address
R$*<$=W.arpa>$*      $#newmail$@2$:address
```

```
R$*<$=W.uucp>$*      $#newmail$@$2$:address
R$*<$=W.$N>$*         $#newmail$@$2$:address
```

(where *address* is the form of the address *newmailer* expects.) Then reinstall the configuration file.

## 8.2. Sending UUCP Mail to A New Host Over the Internet (Gateway Configuration File)

These changes allow you to route mail to internet hosts over the internet, even when the sender has asked that mail be sent via *uucp*. It is strongly recommended that this be done only when *uucp* uses the internet as the communications medium; this ensures the site is really on the internet. Needless to say, this will not work unless the site accepts SMTP connections.

Suppose RIACS talks to the site "faraway.sub.dom" (with *uucp* name *faraway*) using *uucp* running over the internet. To route mail through SMTP but in such a way everyone else thinks the mail was sent through the *uucp* system, add the line

```
R$*<@faraway.uucp>$*      $#utcp$@faraway.sub.dom$:<$1@faraway.sub.dom$2>
```

to ruleset 0, in the section "...handle special hosts". It can go anywhere in that section.

See the section entitled **Mailers** for a description of the *utcp* mailer.

## 8.3. Sending Internet Mail to A New Host Over UUCP (Gateway Configuration File)

These changes allow you to route mail to *uucp* hosts which use the internet domain naming scheme, even when the sender has asked that mail be sent via *uucp*. This simply must be done; it should be undone when (and if) the site joins the internet (in which case perhaps mail routed through *uucp* should be sent via the internet, as described in the section above.)

Suppose RIACS talks to the site "faraway.sub.dom" (with *uucp* name *faraway*) using *uucp* running over a telephone line. It is not possible to pretend the message was sent over the internet, due to the limits of the *uucp* mailers, but adding the line

```
R$*<@faraway.sub.dom>$*      $#uucp$@faraway$:$1$2
```

to ruleset 0, in the section "...handle special hosts" will route such mail over *uucp*. (This rule can go anywhere in that section.) Be aware the address will be run through the *uucp* mailer's sender and recipient address rewriting rules, so the receiving host should get a path to which it can reply. The emphasis is on the word "should."

## 8.4. Adding a New Domain Relay Site (Gateway Configuration File)

Suppose RIACS requires access to the domain "newdom", but that domain can only be reached through the host "relay.dom.net". The goal is to accept addresses like "*user@site.newdom*" and route the mail automatically.

Ruleset 8 is used to handle these cases. Two lines must be added, one for route addressing, and the other for unrouted addressing. In the section of this ruleset entitled "output fake domain stuff in user%host.fake@relay-host syntax", add a rule of the form

```
R$*<@$+.newdom>      $@$1$2.newdom<@relay.dom.net>
```

(which maps "*user<@site.newdom>*" to "*user%site.newdom<@relay.dom.net>*") and to the section of this ruleset entitled "output fake domain stuff in route-specific syntax", add a rule of the form

```
R<@$+.newdom>:$*      $@<@relay.dom.net>:@$1.newdom:$2
```

(which maps "<@site.newdom>:*user@site2*" to "<@relay.dom.net>:@site.newdom:*user@site2*".) Then addresses without the relay site named explicitly will be sent to the relay site anyway.



### 8.5. Change the Gateway Host (Gateway and Non-Gateway Configuration Files)

To do this, both the gateway and non-gateway files must be changed. In the gateway file, change the definition of the class **G** to be all the names of the new gateway host. In the non-gateway file, change the definition of the macro **G** to be the name of the new gateway host (it is recommended this be the fully qualified internet name.) Then install the gateway configuration file on the new gateway host, and the non-gateway configuration file on all other hosts.

This assumes the gateway is to handle all outgoing traffic. If not, a third configuration file must be written, for the UUCP host. See the introductory comments to this section, above. (Translation: good luck!)

### 9. Conclusion

The goal of this work has been to make the RIACS mail system into a low maintenance system and, when maintenance is necessary, easy to maintain. This document is an integral part of that plan, because it is intended to allow a systems programmer to be able to understand how the mail configuration files work, and why the mail system is set up as it is. Needless to say, only time will tell if these rather ambitious goals have been met.

**Acknowledgements:** My thanks to Barry Leiner for his comments on an earlier draft.

### References

- [ALLM84a] Allman, Eric, "Sendmail - An Internetwork Mail Router", in *UNIX System Manager's Manual, 4.2 Berkeley Software Distribution, Virtual VAX-11 Version*, Computer Science Division, Department of Electrical Engineering and Computer Science, University of California, Berkeley, CA (Mar. 1984), as reprinted by the USENIX Association
- [ALLM84b] Allman, Eric, "Sendmail - Installation and Operation Guide Version 4.2", in *UNIX System Manager's Manual, 4.2 Berkeley Software Distribution, Virtual VAX-11 Version*, Computer Science Division, Department of Electrical Engineering and Computer Science, University of California, Berkeley, CA (Mar. 1984), as reprinted by the USENIX Association
- [GILL86] Gilliam, Jeff, USENET message <2821@voder.UUCP>, National Semiconductor, Santa Clara, CA (Dec. 11, 1986).
- [LOVS86] Lovstrand, Lennart, USENET message <377@majestix.liuida.uucp>, Computer Science Department, University of Linköping, Sweden (Dec. 28, 1986).
- [RFC733] Crocker, D. H., Vittal, J. J., Pogran, K. T., and Henderson, D. A., *Standard for the Format of ARPA Network Text Message*, ARPANET Request for Comments No. 733, Network Information Center, SRI International, Menlo Park, CA (Nov. 1977)
- [RFC821] Postel, Jonathan B., *Simple Mail Transfer Protocol*, ARPANET Request for Comments No. 821, Network Information Center SRI International, Menlo Park, CA (Aug. 1982)
- [RFC822] Crocker, David H., *Standard for the Format of ARPA Network Text Message* (revised), ARPANET Request for Comments No. 822, Network Information Center SRI International, Menlo Park, CA (Aug. 1982)
- [RFC920] Postel, J., and Reynolds, J., *Domain Requirements*, ARPANET Request for Comments No. 920, Network Information Center SRI International, Menlo Park, CA (October 1984)
- [ROSE86] Rose, Marshall T., and Romine, John L., *The RAND MH Message Handling System: User's Manual*, UCI Version (Apr. 1986)
- [UPM84] -, *UNIX User's Manual Reference Guide, 4.2 Berkeley Software Distribution, Virtual VAX-11 Version*, Computer Science Division, Department of Electrical Engineering and Computer Science, University of California, Berkeley, CA (Mar. 1984), as

reprinted by the USENIX Association